

AF \$

TRANSMITTAL OF APPEAL BRIEF (Large Entity)

Docket No. **2700**
PQ9-99-159

Re Application Of: **Kubala et al.**

| Serial No. | Filing Date | Examiner | Group Art Unit |
|------------|-------------|----------------|----------------|
| 09/408,470 | 9/28/99 | George L. Opie | 2126 |

Invention: **DYNAMICALLY REDISTRIBUTING SHAREABLE RESOURCES OF A COMPUTING MANAGE THE WORKLOAD OF THAT ENVIRONMENT**

RECEIVED

FEB 09 2004

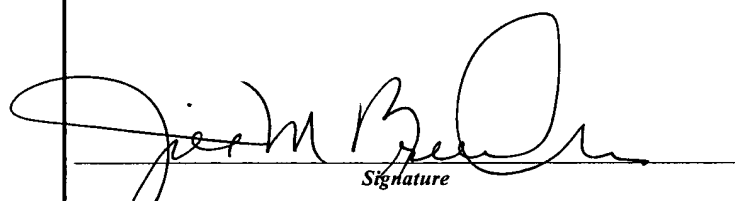
Technology Center 2100

TO THE COMMISSIONER FOR PATENTS:

Transmitted herewith in triplicate is the Appeal Brief in this application, with respect to the Notice of Appeal filed on **December 3, 2003**

The fee for filing this Appeal Brief is: **\$330.00**

- ☐ A check in the amount of the fee is enclosed.
- ☐ The Director has already been authorized to charge fees in this application to a Deposit Account.
- ☒ The Director is hereby authorized to charge any fees which may be required, or credit any overpayment to Deposit Account No. **09-0463 (IBM)**


Signature

Dated: **February 4, 2004**

Jill M. Breedlove, Esq.
Registration No. 32,684
Heslin Rothenberg Farley & Mesiti P.C.
5 Columbia Circle
Albany, New York 12203
Telephone: (518) 452-5600
Facsimile: (518) 452-5579

| |
|---|
| I certify that this document and fee is being deposited on 2/4/04 with the U.S. Postal Service as first class mail under 37 C.F.R. 1.8 and is addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450. |
| <i>Signature of Person Mailing Correspondence</i> |
| Jill M. Breedlove |
| <i>Typed or Printed Name of Person Mailing Correspondence</i> |

cc:



#15/Appel/Brief
T. McBeth Brown
2/11/04

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Appellants: Kubala et al. : Group Art Unit: 2126
Serial No.: 09/408,470 : Examiner: George L. Opie
Filed: 09/28/99 : Appeal No.:

Title: DYNAMICALLY REDISTRIBUTING SHAREABLE RESOURCES OF A
COMPUTING ENVIRONMENT TO MANAGE THE WORKLOAD OF THAT
ENVIRONMENT

RECEIVED

FEB 09 2004

Certificate of Mailing

Technology Center 2100

I hereby certify that this correspondence is being deposited with the U.S. Postal Service as first class mail in an envelope addressed to: Mail Stop Appeal Briefs – Patents, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on February 4, 2004.

Jill M. Breedlove
Attorney for Appellants
Registration No.: 32,684

Date of Signature: February 4, 2004.

Mail Stop Appeal Briefs – Patents
Commissioner for Patents
P.O. Box 1450, Alexandria, VA 22313-1450

02/09/2004 AWONDAF1 00000090 090463 09408470

01 FC:1402 330.00 DA

Brief of Appellants

Dear Sir:

This is an appeal from a final rejection, dated September 3, 2003 rejecting claims 1-65, all the claims being considered in the above-identified application. This Brief is accompanied by a transmittal letter authorizing the charging of appellants' deposit account for payment of the requisite fee set forth in 37 C.F.R. §1.17(c).

Real Party in Interest

This application is assigned to **International Business Machines Corporation** by virtue of an assignment executed by the co-inventors on November 30, 1999 and December 3, 1999 and November 19, 1999; and recorded with the United States Patent and Trademark Office at reel <number>, frame <number>, on <month> <date>, <yr>. Therefore, the real party in interest is **International Business Machines Corporation**.

Related Appeals and Interferences

To the knowledge of the appellants, appellants' undersigned legal representative, and the assignee, there are no other appeals or interferences which will directly affect or be directly affected by or have a bearing on the Board's decision in the instant appeal.

Status of Claims

This patent application was filed on September 28, 1999 with the United States Patent and Trademark Office. As filed, the application included eight (8) independent claims (i.e., claims 1, 13, 22, 34, 43, 44, 45 & 57).

In an initial Office Action dated February 19, 2003 claims 1-65 were rejected under 35 U.S.C. 103(a) as being unpatentable over Admitted Prior Art (APA) from the BACKGROUND ART section of the application in view of Maeurer et al. (U.S. Patent 5,301,323). In appellants' response dated June 19, 2003 (with which a request for a one-month extension of time and requisite fee were enclosed), claims 1, 13, 15, 18-20, 22, 34, 36, 39-41, 43-45, 57, 59, 62-64 were amended.

In a second and final Office Action dated September 3, 2003, claims 1-65 were rejected under 35 U.S.C. 103(a) as being unpatentable over Admitted Prior Art (APA) from the BACKGROUND ART in view of Maeurer et al. (U.S. Patent 5,301,323) and in view of Onodera et al. (U.S. Patent 5,996,026). The Office Action rendered the new ground of rejection as

necessitated by the appellants' amendment of June 19, 2003. In appellants' response dated November 3, 2003, appellants proffered remarks in Response to the final Office Action; no claims were amended.

A Notice of Appeal to the Board of Patent Appeals and Interferences was filed on December 3, 2003. The status of the claims is therefore as follows:

Claims allowed – none;

Claims objected to – none;

Claims rejected – 1-65; and

Claims canceled – none.

Appellants are appealing the rejection of claims 1-65.

Status of Amendments

The claims as set forth in the Appendix include all prior entered claim amendments.

Summary of Invention

The present invention is directed to the problem of managing workload in a computing environment.

More particularly, the present invention is a method (claim 1), system (claims 22 and 43), and program storage device embodying a program of instructions readable and executable by a machine to perform the method (claim 43) for managing workload across two or more partitions of a plurality of partitions of a computing environment, wherein a partition has one or more central processors allocated thereto, the managing comprising dynamically adjusting allocation of a shareable resource, wherein workload goals are being balanced. (see FIG. 1A, specification page 11, lines 1-23, through page 12, lines 1-12). Shareable resources may include CPU resources, logical processor resources, input/output resources, channel resources, coprocessors,

network adapters, and memory, among others (see specification page 15, lines 6-16; claims 3, 14, 24, 35, 47 and 58).

In one aspect, the method comprises dynamically adjusting allocation of a shareable resource based on priority (FIG. 9; specification page 16, lines 12-15; page 40, lines 20-30, through page 47; claims 5, 19, 26, 40, 49 and 63).

In a further aspect, the dynamically adjusting comprises moving the shareable resources to the physical work, i.e., moving at least a portion of the shareable resource from one partition to another partition (see specification page 17, lines 19-29; page 26, lines 20-30, through page 27, lines 1-2; and claims 4, 18, 25, 39, 48 and 62).

In a still further aspect, the present invention is a method (claim 13), system (claims 34 and 44), and article of manufacture having computer readable code means for performing the method (claim 57), wherein the computing environment comprises two or more logical partitions concurrently sharing at least one shareable resource (see FIGs. 1-3, specification page 10, lines 6-9; and claims 13, 34, 44 and 57).

Issue

1. Whether claims 1-65 were rendered obvious under 35 U.S.C. 103(a) by Admitted Prior Art (APA) in view of Maeurer et al. (U.S. Patent 5,301,323) and in view of Onodera et al. (U.S. Patent 5,996,026).

Grouping of Claims

As to the rejections applied against claims 1-65, it is appellants' intention that the rejected claims do not stand or fall together. For example, appellants' arguments provide a grouping of claims, the following principal groups of claims being included herein as follows:

- I. Claims 1-3, 5-17, 19-24, 26-38, 40-47, 49-61 and 63-65; and
- II. Claims 4, 18, 25, 39, 48, and 62.

As understood, the claims of one group of claims do not stand or fall with any other group of claims. Further, appellants' respectfully submit that the claims within each group do not stand or fall together.

Argument

Group I: Claims 1-3, 5-17, 19-24, 26-38, 40-47, 49-61 and 63-65

As noted, claims 1-3, 5-17, 19-24, 26-38, 40-47, 49-61 and 63-65 stand rejected under 35 U.S.C. §103(a) as being obvious over Admitted Prior Art (APA) in view of Maurer et al. (U.S. Patent 5,301,323) and in view of Onodera et al. (U.S. Patent 5,996,026). Appellants respectfully request reversal of this rejection.

At the outset, the appellants respectfully traverse the suggested combination of APA, Maurer et al., and Onodera et al. to the extent that it is alleged that the combination suggests their claimed invention. The justification given in the final Office Action for the combination of APA and Maurer et al. is "because the dynamic reallocation/reconfiguration of a device with respect to channel/partitions provides a flexibility [sic] for assigning resources responsive to system load conditions." The Onodera et al. patent is then cited because "[t]he APA/Maurer [combination] does not explicitly disclose the partition having one or more central processors allocated thereto and "because the allocation of 'one' or more processors to a partition would facilitate processing by the partition."

Although the appellants respectfully submit that the Office Actions mischaracterize to some extent the APA (as will be explained hereinbelow), the appellants agree that the APA describes a logically partitioned system and management of workload within the partitions of the system.

Maurer et al., however, do not teach or suggest partitions, as alleged in the Office Actions. Maurer et al. teach a data processing system having a channel path management program that is periodically executed to monitor utilization of control units and subchannels. If

the degree of utilization exceeds a predetermined threshold, then the channel path configuration is changed. Maurer et al. define “configuration” to mean the physical connections in the system as well as channel subsystem and software settings (Maurer et al., column 2, lines 7-12). The data processing system of Maurer et al. has one CPU (see, e.g., FIG. 1 and column 4, lines 20-22).

The appellants respectfully submit that the justification given in the Office Actions, and set forth hereinabove, does not identify an adequate teaching, suggestion or incentive in the art itself for the suggested combination of APA and Maurer et al. In particular, as agreed, the APA describes a logically partitioned system and management of workload within the partitions of the system. Maurer et al. do not teach or suggest a partitioned system. Thus, it is believed that any such combination would only be drawn from appellants’ own disclosure.

The appellants respectfully submit, therefore, that the only suggestion or incentive for combining the two teachings in the manner set forth in the Office Actions is presented in the appellants’ own disclosure, which according to established principles, cannot be used as a reference against their claimed invention.

In order to justify the suggested combination of APA and Maurer et al. without reference to the appellants’ disclosure, the Office Actions equate partitions with channel paths. The appellants respectfully, but vigorously, disagree. Maurer et al. define channel paths in column 1, lines 40-45, as follows:

The connections between the channel subsystem and the control units are referred to as channel paths. The channel paths may extend through and be routed by switches. An ‘I/O configuration’ comprises channel paths and control units, I/O devices and their connections including any switches.

Partitions are not channel paths. More particularly, the appellants set forth a definition for logical partitioning on page 3, lines 10-18, of their specification as follows:

Logical partitioning allows the establishment of a plurality of system images within a single physical machine or central processor complex (CPC). Each system image is capable of operating as if it was a separate computer system. That is, each logical partition can be independently reset, initially loaded with an operation system that may be different for each logical partition, and operate with different software programs using different input/output (I/O) devices.

Hence, from a careful reading of APA and Maeurer et al., partitions are not equivalent to channel paths, such that the justification of the combination of APA and Maurer et al. does not follow. Moreover, Onodera et al. does not make up for the deficiencies of either of these references to render obvious the appellants' claimed invention. In particular, Onodera et al. is cited because "the APA/Maurer [combination] does not explicitly disclose the partition having one or more central processors allocated thereto." The appellants agree that Onodera et al. describe using logical partitions. Further, Onodera et al. describe a configuration method for a hypervisor, which virtualizes a channel path configuration that exceeds a predetermined threshold (e.g., 256 channel paths) to a channel path configuration that can be utilized (e.g., less than or equal to 256 channel paths). However, there is nothing in Onodera about dynamically managing shared resources, as recited by the appellants, i.e., by moving shared resources to the work as needed. Thus, merely by using logical partitions, which the appellants have agreed is not new, Onodera et al. do not overcome the deficiencies of APA and Maurer et al., described hereinabove.

Therefore, neither APA nor Maeurer et al. nor Onodera et al. teaches or suggests managing workload across two or more partitions of a plurality of partitions of a computing environment, wherein the managing comprises dynamically adjusting allocation of a shareable resource of at least one partition, as recited by the appellants in independent claims 1, 13, 22, 34, 43, 44, 45 and 57, and the claims dependent therefrom. Hence, the appellants' invention, as recited in the aforementioned claims, is not rendered obvious by the suggested combination of APA, Maeurer et al. and Onodera et al. Appellants thus respectfully request reversal of the rejection of these claims.

Group II: Claims 4, 18, 25, 39, 48, and 62

As noted, claims 4, 18, 25, 39, 48, and 62 stand rejected under 35 U.S.C. §103(a) as being obvious over Admitted Prior Art (APA) in view of Maeurer et al. (U.S. Patent 5,301,323) and in view of Onodera et al. (U.S. Patent 5,996,026). Appellants respectfully request reversal of this rejection.

These dependent claims are believed to be allowable for the same reasons as the independent claims of Group I, as well as for their own additional characterizations. In this regard, appellants respectfully traverse the secondary obviousness rejections to claims 4, 18, 25, 39, 48, and 62 based on APA in view of Maeurer et al. and in view of Onodera et al. et al. for the same reasons noted above with respect to the claims of Group I.

The dependent claims of Group II further recite that dynamically adjusting allocation of a shareable resource comprises “moving at least a portion of said shareable resource from one partition to at least one other partition.” That is, the appellants recite moving shareable resources to the work.

The final Office Action relies on APA as “teach[ing] a method of managing workload of a computing environment (moving work to the physical resources of the system) [,] said method comprising: managing workload across two or more partitions of a plurality of partitions of said computing environment (S/390... workload managers are used to manage the workload within and among the partitions.” The appellants respectfully submit that this is a mischaracterization of the APA. In particular, the specification, page 4, lines 9-21, states as follows:

One important aspect of a logically partitioned system is the management of workload running within the partitions of that system. In S/390 systems, for example, workload managers are used to manage the workload within and among the partitions. The workload managers attempt to balance the workload of the partitions *by moving work to the physical resources of the system.*” (emphasis added).

Thus, contrary to the characterization in the final Office Action, the APA explicitly refers to workload managers that move work to physical resources to the system. This is in contrast to the appellants' recitation in claims 4, 18, 25, 39, 48 and 62 that "said dynamically adjusting comprises moving at least a portion of said shareable resource from one partition to at least one other partition," i.e., moving resources to the work.

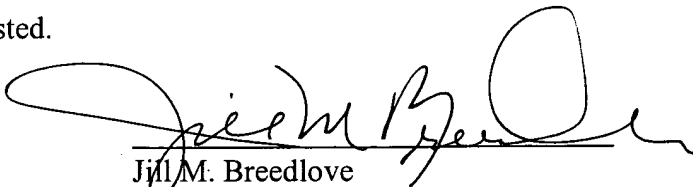
Therefore, the APA does not teach moving physical resources to the work. Moreover, in addition to the other deficiencies of the APA, Maeurer et al. and Onodera et al. described hereinabove with respect to Group I, none of these references teaches or suggests moving physical resources to the work.

Claims 4, 18, 25, 39, 48 and 62 are thus believed to be patentable to the appellants over these references under 35 USC 103(a). Appellants thus respectfully request reversal of the rejection to these claims.

Conclusion

Appellants respectfully request reversal of the rejections as set forth in the final Office Action. Appellants submit that their claimed invention would not have been rendered obvious by Admitted Prior Art, Maeurer et al. and Onodera et al. These references do not individually or in combination teach, suggest, or imply appellants' recited structure, which includes, for example, managing workload across two or more partitions of a plurality of partitions of a computing environment, wherein the managing comprises dynamically adjusting allocation of a shareable resource, and wherein workload goals are being balanced.

For all of the above reasons, appellants allege error in rejecting their claims as being obvious based on Admitted Prior Art, Maeurer et al. and Onodera et al. Accordingly, reversal of the rejections is respectfully requested.

A handwritten signature in black ink, appearing to read "Jill M. Breedlove", written over a horizontal line.

Jill M. Breedlove
Attorney for Appellants
Registration No.: 32,684

Dated: February 4, 2004.

HESLIN ROTHENBERG FARLEY & MESITI P.C.

5 Columbia Circle

Albany, New York 12203-5160

Telephone: (518) 452-5600

Facsimile: (518) 452-5579

Appendix A

1. A method of managing workload of a computing environment, said method comprising:

managing workload across two or more partitions of a plurality of partitions of said computing environment, wherein a partition has one or more central processors allocated thereto;

said managing comprising dynamically adjusting allocation of a shareable resource of at least one partition of said two or more partitions, wherein workload goals of said two or more partitions are being balanced.

2. The method of claim 1, wherein said dynamically adjusting is performed transparently to work processing within said at one least one partition.

3. The method of claim 1, wherein said shareable resource comprises at least one of central processing unit resources, logical processor resources, input/output resources, channel resources, coprocessors, network adapters, and memory.

4. The method of claim 1, wherein said dynamically adjusting comprises moving at least a portion of said shareable resource from one partition to at least one other partition.

5. The method of claim 1, wherein said dynamically adjusting comprises managing said shareable resource among said two or more partitions based on priority.

6. The method of claim 1, wherein said dynamically adjusting comprises assigning said shareable resource among said two or more partitions based on percentage allocation, wherein each partition of said two or more partitions is assigned a percentage of said shareable resource.

7. The method of claim 1, wherein said partitions are logical partitions.

8. The method of claim 1, wherein said dynamically adjusting comprises adjusting allocation of a plurality of shareable resources.

9. The method of claim 1, wherein said dynamically adjusting is controlled at least in part by at least one workload manager of said computing environment.

10. The method of claim 1, wherein said dynamically adjusting comprises increasing allocation of said shareable resource.

11. The method of claim 1, wherein said dynamically adjusting comprises decreasing allocation of said shareable resource.

12. The method of claim 1, wherein said dynamically adjusting is performed without a requirement for data sharing.

13. A method of managing workload of a computing environment, said method comprising:

managing workload across two or more logical partitions of a plurality of logical partitions of said computing environment, wherein said two or more logical partitions concurrently share at least one shareable resource;

said managing comprising dynamically adjusting allocation of said shareable resource of at least one logical partition of said two or more logical partitions.

14. The method of claim 13, wherein said shareable resource comprises at least one of central processing unit resources, logical processor resources, input/output resources, channel resources, coprocessors, network adapters, and memory.

15. The method of claim 13, wherein workload goals of said two or more logical partitions are being balanced.

16. The method of claim 13, wherein said dynamically adjusting comprises increasing allocation of said shareable resource.

17. The method of claim 13, wherein said dynamically adjusting comprises decreasing allocation of said shareable resource.

18. The method of claim 13, wherein said dynamically adjusting comprises moving at least a portion of said shareable resource from one logical partition to at least one other logical partition.

19. The method of claim 13, wherein said dynamically adjusting comprises managing said shareable resource among said two or more logical partitions based on priority.

20. The method of claim 13, wherein said dynamically adjusting comprises assigning said shareable resource among said two or more logical partitions based on percentage allocation, wherein each logical partition of said two or more logical partitions is assigned a percentage of said shareable resource.

21. The method of claim 13, wherein said dynamically adjusting comprises adjusting allocation of a plurality of shareable resources.

22. A system of managing workload of a computing environment, said system comprising:

means for managing workload across two or more partitions of a plurality of partitions of said computing environment, wherein a partition has one or more central processors allocated thereto;

said means for managing comprising means for dynamically adjusting allocation of a shareable resource of at least one partition of said two or more partitions, wherein workload goals of said two or more partitions are being balanced.

23. The system of claim 22, wherein the dynamically adjusting is performed transparently to work processing within said at one least one partition.

24. The system of claim 22, wherein said shareable resource comprises at least one of central processing unit resources, logical processor resources, input/output resources, channel resources, coprocessors, network adapters, and memory.

25. The system of claim 22, wherein said means for dynamically adjusting comprises means for moving at least a portion of said shareable resource from one partition to at least one other partition.

26. The system of claim 22, wherein said means for dynamically adjusting comprises means for managing said shareable resource among said two or more partitions based on priority.

27. The system of claim 22, wherein said means for dynamically adjusting comprises means for assigning said shareable resource among said two or more partitions based on percentage allocation, wherein each partition of said two or more partitions is assigned a percentage of said shareable resource.

28. The system of claim 22, wherein said partitions are logical partitions.
29. The system of claim 22, wherein said means for dynamically adjusting comprises means for adjusting allocation of a plurality of shareable resources.
30. The system of claim 22, wherein said means for dynamically adjusting is controlled at least in part by at least one workload manager of said computing environment.
31. The system of claim 22, wherein said means for dynamically adjusting comprises means for increasing allocation of said shareable resource.
32. The system of claim 22, wherein said means for dynamically adjusting comprises means for decreasing allocation of said shareable resource.
33. The system of claim 22, wherein said means for dynamically adjusting is performed without a requirement for data sharing.
34. A system of managing workload of a computing environment, said system comprising:
- means for managing workload across two or more logical partitions of a plurality of logical partitions of said computing environment, wherein said two or more logical partitions concurrently share at least one shareable resource;
- said means for managing comprising means for dynamically adjusting allocation of said shareable resource of at least one logical partition of said two or more logical partitions.

35. The system of claim 34, wherein said shareable resource comprises at least one of central processing unit resources, logical processor resources, input/output resources, channel resources, coprocessors, network adapters, and memory.

36. The system of claim 34, wherein workload goals of said two or more logical partitions are being balanced.

37. The system of claim 34, wherein said means for dynamically adjusting comprises increasing allocation of said shareable resource.

38. The system of claim 34, wherein said means for dynamically adjusting comprises means for decreasing allocation of said shareable resource.

39. The system of claim 34, wherein said means for dynamically adjusting comprises means for moving at least a portion of said shareable resource from one logical partition to at least one other logical partition.

40. The system of claim 34, wherein said means for dynamically adjusting comprises means for managing said shareable resource among said two or more logical partitions based on priority.

41. The system of claim 34, wherein said means for dynamically adjusting comprises means for assigning said shareable resource among said two or more logical partitions based on percentage allocation, wherein each logical partition of said two or more logical partitions is assigned a percentage of said shareable resource.

42. The system of claim 34, wherein said means for dynamically adjusting comprises means for adjusting allocation of a plurality of shareable resources.

43. A system of managing workload of a computing environment, said system comprising:

a processor adapted to manage workload across two or more partitions of a plurality of partitions of said computing environment, wherein a partition has one or more central processors allocated thereto; and

wherein said managing comprises dynamically adjusting allocation of a shareable resource of at least one partition of said two or more partitions, wherein workload goals of said two or more partitions are being balanced.

44. A system of managing workload of a computing environment, said system comprising:

a processor adapted to manage workload across two or more logical partitions of a plurality of logical partitions of said computing environment, wherein said two or more logical partitions concurrently share at least one shareable resource; and

wherein said managing comprises dynamically adjusting allocation of said shareable resource of at least one logical partition of said two or more logical partitions.

45. At least one program storage device readable by a machine, tangibly embodying at least one program of instructions executable by the machine to perform a method of managing workload of a computing environment, said method comprising:

managing workload across two or more partitions of a plurality of partitions of said computing environment, wherein a partition has one or more central processors allocated thereto;

said managing comprising dynamically adjusting allocation of a shareable resource of at least one partition of said two or more partitions, wherein workload goals of said two or more partitions are being balanced.

46. The at least one program storage device of claim 45, wherein said adjusting is performed transparently to work processing within said at one least one partition.

47. The at least one program storage device of claim 45, wherein said shareable resource comprises at least one of central processing unit resources, logical processor resources, input/output resources, channel resources, coprocessors, network adapters, and memory.

48. The at least one program storage device of claim 45, wherein said dynamically adjusting comprises moving at least a portion of said shareable resource from one partition to at least one other partition.

49. The at least one program storage device of claim 45, wherein said dynamically adjusting comprises managing said shareable resource among said two or more partitions based on priority.

50. The at least one program storage device of claim 45, wherein said dynamically adjusting comprises assigning said shareable resource among said two or more partitions based on percentage allocation, wherein each partition of said two or more partitions is assigned a percentage of said shareable resource.

51. The at least one program storage device of claim 45, wherein said partitions are logical partitions.

52. The at least one program storage device of claim 45, wherein said dynamically adjusting comprises adjusting allocation of a plurality of shareable resources.

53. The at least one program storage device of claim 45, wherein said dynamically adjusting is controlled at least in part by at least one workload manager of said computing environment.

54. The at least one program storage device of claim 45, wherein said dynamically adjusting comprises increasing allocation of said shareable resource.

55. The at least one program storage device of claim 45, wherein said dynamically adjusting comprises decreasing allocation of said shareable resource.

56. The at least one program storage device of claim 45, wherein said dynamically adjusting is performed without a requirement for data sharing.

57. An article of manufacture, comprising:

at least one computer usable medium having computer readable program code means embodied therein for causing the managing of workload of a computing environment, the computer readable program code means in said article of manufacture comprising:

computer readable program code means for causing a computer to manage workload across two or more logical partitions of a plurality of logical partitions of said computing environment, wherein said two or more logical partitions concurrently share at least one shareable resource;

said computer readable program code means for causing a computer to manage comprising computer readable program code means for causing a computer to dynamically adjust allocation of said shareable resource of at least one logical partition of said two or more logical partitions.

58. The article of manufacture of claim 57, wherein said shareable resource comprises at least one of central processing unit resources, logical processor resources, input/output resources, channel resources, coprocessors, network adapters, and memory.

59. The article of manufacture of claim 57, wherein workload goals of said two or more logical partitions are being balanced.

60. The article of manufacture of claim 57, wherein said computer readable program code means for causing a computer to dynamically adjust comprises computer readable program code means for causing a computer to increase allocation of said shareable resource.

61. The article of manufacture of claim 57, wherein said computer readable program code means for causing a computer to dynamically adjust comprises computer readable program code means for causing a computer to decrease allocation of said shareable resource.

62. The article of manufacture of claim 57, wherein said computer readable program code means for causing a computer to dynamically adjust comprises computer readable program code means for causing a computer to move at least a portion of said shareable resource from one logical partition to at least one other logical partition.

63. The article of manufacture of claim 57, wherein said computer readable program code means for causing a computer to dynamically adjust comprises computer readable program code means for causing a computer to manage said shareable resource among said two or more logical partitions based on priority.

64. The article of manufacture of claim 57, wherein said computer readable program code means for causing a computer to dynamically adjust comprises computer readable program code means for causing a computer to assign said shareable resource among said two or more logical partitions based on percentage allocation, wherein each logical partition of said two or more logical partitions is assigned a percentage of said shareable resource.

65. The article of manufacture of claim 57, wherein said computer readable program code means for causing a computer to dynamically adjust comprises computer readable program code means for causing a computer to adjust allocation of a plurality of shareable resources.

* * * * *